

GAMEPLAYQA: A Benchmarking Framework for Decision-Dense POV-Synced Multi-Video Understanding of 3D Virtual Agents

Yunzhe Wang, Runhui Xu, Kexin Zheng, Tianyi Zhang,
Jayavibhav Niranjana Kogundi, Soham Hans, Volkan Ustun

University of Southern California

{yunzhewa, runhuixu, kexinzhe, tzhang62, jniranja, sohamhan, ustun}@usc.edu

<https://hats-ict.github.io/gameplayqa/>

Abstract

Multimodal LLMs are increasingly deployed as perceptual backbones for autonomous agents in 3D environments, from robotics to virtual worlds. These applications require agents to perceive rapid state changes, attribute actions to the correct entities, and reason about concurrent multi-agent behaviors from a first-person perspective, capabilities that existing benchmarks do not adequately evaluate. We introduce GAMEPLAYQA, a framework for evaluating agentic-centric perception and reasoning through video understanding. Specifically, we densely annotate multiplayer 3D gameplay videos at 1.22 labels/second, with time-synced, concurrent captions of states, actions, and events structured around a triadic system of **Self**, **Other Agents**, and the **World**, a natural decomposition for multi-agent environments. From these annotations, we refined 2.4K diagnostic QA pairs organized into three levels of cognitive complexity, accompanied by a structured distractor taxonomy that enables fine-grained analysis of where models hallucinate. Evaluation of frontier MLLMs reveals a substantial gap from human performance, with common failures in temporal and cross-video grounding, agent-role attribution, and handling the decision density of the game. We hope GAMEPLAYQA stimulates future research at the intersection of embodied AI, agentic perception, and world modeling.

1 Introduction

Recent advances in Multimodal Large Language Models (MLLMs) have demonstrated remarkable capabilities in advanced reasoning, multimodality, and agency (Comanici et al., 2025; OpenAI et al., 2024; Anthropic, 2025; Bai et al., 2025), positioning them as promising decision-making backbones for autonomous agents in Robotics (Zitkovich et al., 2023; Gemini Robotics Team et al., 2025), Computer Use (He et al., 2024; Zhang et al., 2026) and 3D virtual agents (SIMA Team et al., 2024;

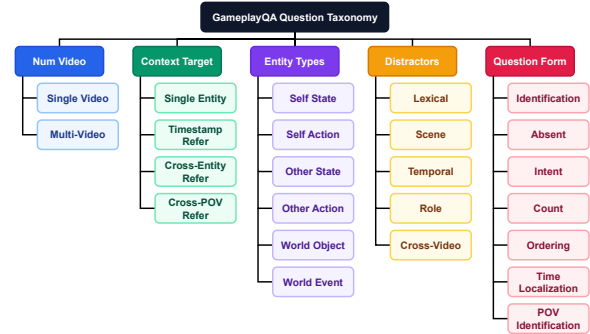


Figure 1: Question taxonomy of GAMEPLAYQA. Questions are organized along two axes: *Entity* (Self, Other, World) and *Temporal Nature* (Action/State for agents, Object/Event for world), yielding six primitive label types. These primitives compose into 15 task categories across three cognitive levels: single-reference perception (L1), temporal reasoning (L2), and cross-video understanding (L3). See Sec. 3.1 and Table 2 for details.

SIMA team et al., 2025; Yue et al., 2026). These applications require perception capabilities beyond passive scene description. Drawing on perspectives from embodied cognition and multi-agent reasoning (Hernandez-Leal et al., 2019), we identify three core requirements for agentic perception in goal-directed environments: (1) **dense state-action tracking**: capturing rapid transitions in the agent’s own states and actions; (2) **other-agent modeling**: reasoning about the behaviors and intentions of other autonomous entities; and (3) **environment grounding**: tracking persistent and transient elements of the shared world.

However, current video understanding benchmarks are ill-equipped to diagnose these agentic requirements for three primary reasons. First, the majority of existing evaluation sets suffer from a lack of embodiment and agency grounding (Majumdar et al., 2024; Yang et al., 2025; Dang et al., 2025); they are often composed of slow-paced, passive observations that lack the high-frequency state transitions and dense decision-making loops re-

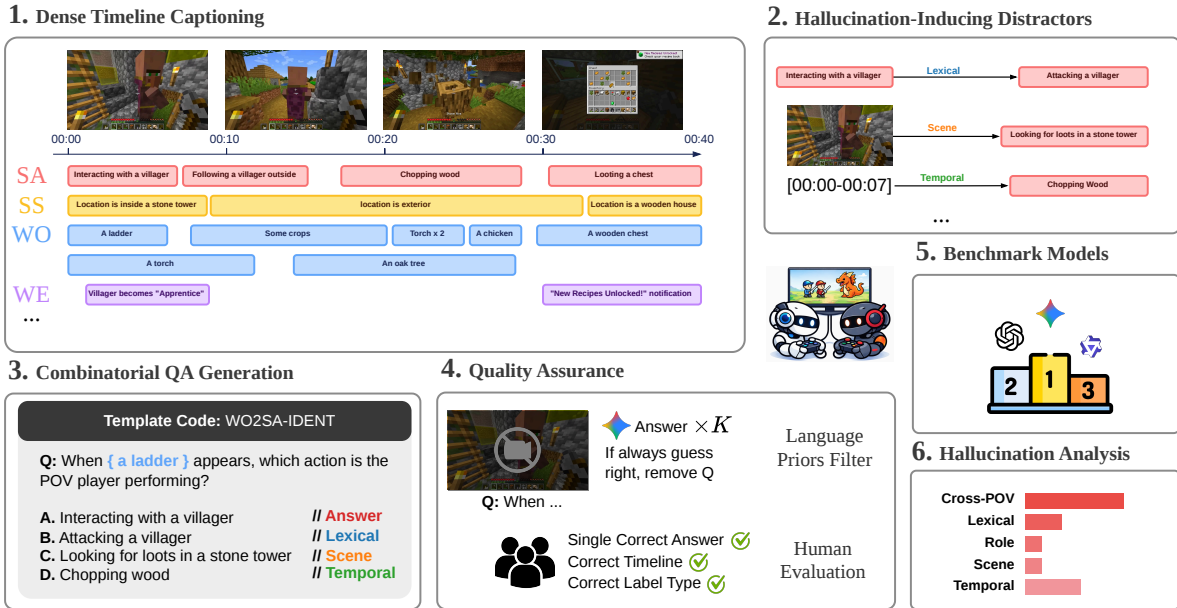


Figure 2: Overview of GAMEPLAYQA. Gameplay videos undergo (1) dense multi-track temporal captioning on 6 types of target entities (Sec. 3.2), (2) captioning includes negative labels for hallucination-inducing distractors, and (3) QA pairs are generated through a combinatorial template-based algorithm (Sec. 3.3). After (4) quality assurance (Sec. 3.4), the benchmark enables (5) model evaluation with (6) fine-grained hallucination analysis (Sec. 4.2).

quired to stress-test a model’s understanding of intentional action. Second, these benchmarks are largely not hallucination-diagnosable, providing global performance metrics while lacking the granular, multi-faceted annotation needed to identify whether a failure stems from temporal misinterpretation, object fabrication, or role confusion (Bai et al., 2024; Seth et al., 2025; Tu et al., 2025). Finally, current protocols exhibit a significant lack of multi-video understanding (Peng et al., 2025), focusing almost exclusively on single-viewpoint perception. Multi-video understanding is important in domains such as *sports analytics* leveraging various camera angles and *autonomous driving* requiring information fusion from multiple surround cameras. In *esports* and gaming, cross-POV synchronization and collective reasoning, skills that are fundamental to interpreting multi-agent collaboration in interactive 3D spaces (Long et al., 2024; Savva et al., 2026), are also crucial.

To bridge this gap, we introduce GAMEPLAYQA, a comprehensive benchmarking framework, not merely a static evaluation set, but an end-to-end pipeline encompassing structured annotation protocols, automated question generation, and diagnosable error analysis, designed to evaluate the cognitive foundations of agency in 3D virtual environments. We utilize 3D gameplay as a high-density “cognitive sandbox” where states and con-

sequences are deterministic and decision-making is fast-paced. We meticulously annotate synchronized gameplay videos from 9 multiplayer commercial games at a decision density of $\rho \approx 1.22$ labels/second (Eq. 1), using a timeline-based dense captioning mechanism structured around a **Self-Other-World** entity decomposition. This tripartite schema combined with the properties of 3D gameplay directly addresses the three core agentic requirements identified above: *Self* captures the POV agent’s own states and actions for dense state-action tracking; *Other* models external agents’ behaviors and intentions; and *World* grounds perception in persistent environmental elements and transient events (Fig. 5).

Leveraging these annotations, we propose a combinatorial template-based algorithm that generates 2.4K QA pairs organized into a multi-faceted taxonomy spanning three cognitive levels: (1) basic perception, (2) temporal reasoning, and (3) cross-video understanding. The algorithm initially produces 400K candidate pairs and we downsample to 4K to enforce balanced category coverage before quality assurance yields the final set. A key innovation is our structured distractor taxonomy: by categorizing incorrect options as lexical, temporal, or role-based confusions, we can systematically diagnose model hallucination through multiple-choice questions. Evaluation of state-of-the-art MLLMs

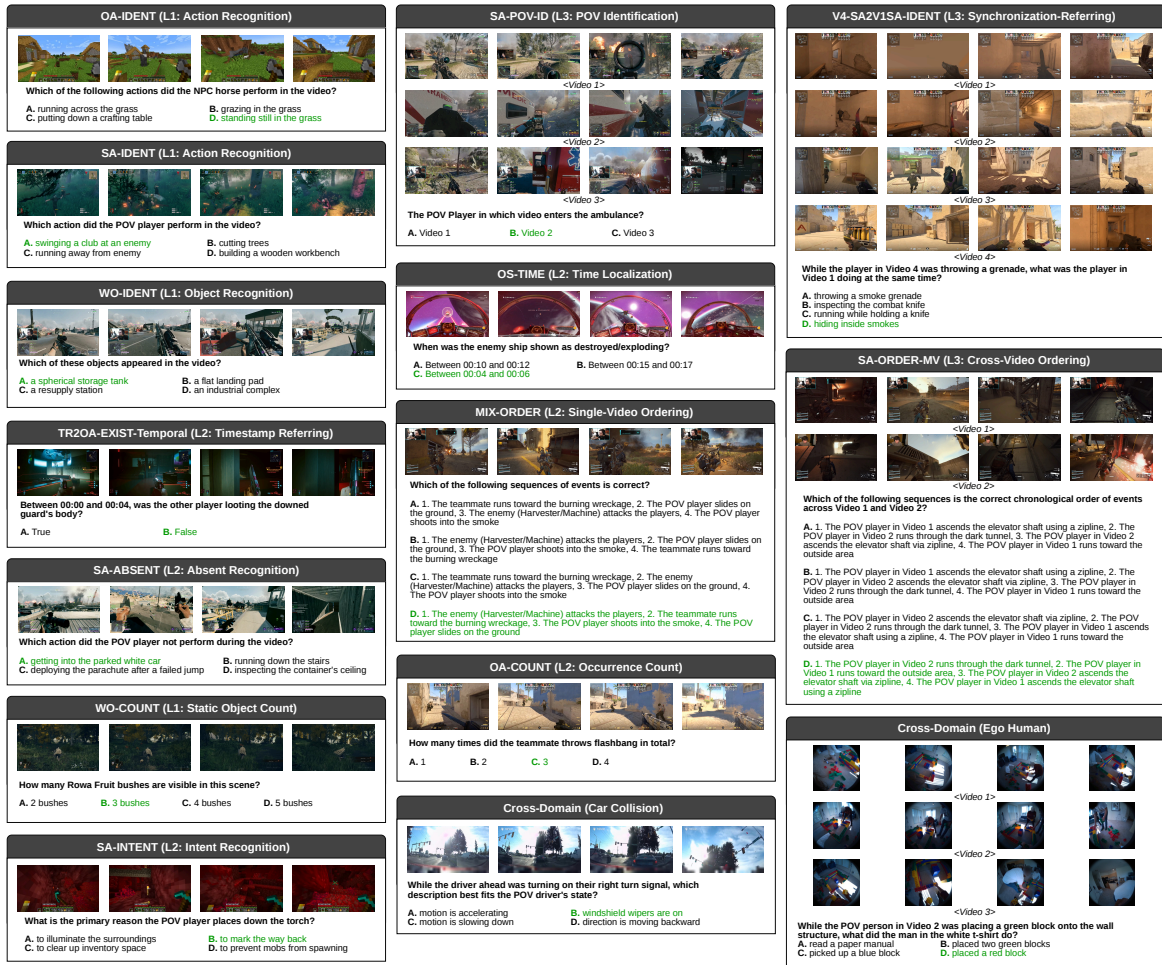


Figure 3: Example questions from GAMEPLAYQA across different question codes and cognitive levels. Each example shows video frames paired with the corresponding QA pair, illustrating the progression from basic perception (L1) to temporal reasoning (L2) to cross-video understanding (L3). Additional cross-domain examples from car collision and egocentric human activity videos demonstrate the generalizability of the framework.

reveals a performance gap against human, with models struggling when: (1) the game is fast-paced and decision-dense, (2) questions concern other agents or entities rather than the egocentric player, and (3) cross-video understanding and temporal grounding over long horizons are required.

In summary, our contributions are threefold:

- We introduce an end-to-end benchmarking framework with structured taxonomy, annotation schema, combinatorial QA generation, and diagnosable error analysis, enabling reproducible evaluation pipelines that can scale to new games and domains.
- We release a benchmark of 2.4K QA pairs from 9 multiplayer games with synchronized multi-POV videos, filling a critical gap in evaluating the dense, multi-agent perception required for embodied AI.

- Benchmarking frontier MLLMs against human evaluation reveals a performance gap, with fine-grained diagnostic analysis through structured distractors revealing that models struggle with fast-paced decision-dense scenarios, other-agent modeling, cross-video synchronization grounding, and temporal reasoning over long horizons.

2 Related Work

Multimodal Large Language Models Recent progress in MLLMs has significantly expanded the ability of AI systems to perceive and reason over visual inputs (Comanici et al., 2025; OpenAI et al., 2024; Anthropic, 2025; Bai et al., 2023). Many recent MLLMs have been proposed to be video-native for video understanding (Cheng et al., 2024; Comanici et al., 2025; Li et al., 2024b). These systems can process extended visual streams; however,

Benchmark	Domain	Agent-Centric	Multi-POV	Diagnostic	Annotation	#Q	#Vid
MVBench (Li et al., 2024a)	General	✗	✗	✓	Auto	4,000	4,000
LongVideoBench (Wu et al., 2024)	General	✗	✗	✓	Human	6,678	3,763
MVU-Eval (Peng et al., 2025)	General	✗	✓	✗	Human, Auto	1,824	4,959
MovieQA (Tapaswi et al., 2016)	Movie	✗	✗	✗	Human	15k	408
TVQA (Lei et al., 2018)	TV Shows	✗	✗	✗	Human	152k	21.8k
MarioQA (Mun et al., 2017)	Gameplay	✓	✗	✗	Auto	188k	13 hrs
PhysGame (Cao et al., 2024)	Game Glitches	✗	✗	✓	Human	880	880
VideoGameQA-Bench (Taesiri et al., 2025)	Game Glitches	✗	✗	✓	Human, Auto	4,786	1.2k
Ego4D (Grauman et al., 2022)	Daily/Ego	✓	✗	✓	Human	74k	3.67k hrs
EgoSchema (Mangalam et al., 2023)	Daily/Ego	✓	✗	✓	Human	5.1k	5.1k
EgoIllusion (Seth et al., 2025)	Daily/Ego	✓	✗	✓	Human, Auto	8k	1.4k
GameplayQA (Ours)	Gameplay	✓	✓	✓	Human, Auto	2,365	100

Table 1: Comparison of relevant video understanding benchmarks.

Context Scope	Task	Description	Example Question Codes	#Q	Dur.
Single Reference (L1, 469)	Action Recognition	Identify or verify existence of self & others’ actions	SA-IDENT, OA-EXIST, ...	162	10.0s
	State Recognition	Identify or verify existence of self & others’ states	SS-IDENT, OS-EXIST, ...	147	10.1s
	Object Recognition	Identify or verify existence of world objects in scene	WO-IDENT, WO-EXIST	70	9.3s
	Event Recognition	Identify world events occurring in the environment	WE-IDENT	61	8.4s
	Static Object Count	Count static objects present in the scene	WO-COUNT	29	21.3s
Temporal (L2, 1383)	Cross-Entity Referring	Link one entity to another (X2Y reasoning)	SA2SS-IDENT, WO2SS-EXIST, ...	423	23.0s
	Timestamp Referring	Given time range [t1-t2], identify what entity exists	TR2SS-IDENT, TR2SA-IDENT, ...	81	24.3s
	Time Localization	Locate exact timestamp when an event occurred	SA-TIME, WE-TIME, ...	281	28.4s
	Absence Recognition	Identify actions/states that did not occur over a timespan	SA-ABSENT, SS-ABSENT, ...	195	38.9s
	Occurrence Count	Count how many times an action/event happened	SA-COUNT, OA-COUNT, WE-COUNT	75	26.4s
	Ordering	Determine temporal order sequence of actions/events	SA-ORDER, OA-ORDER, MIX-ORDER	180	32.6s
	Intent Identification	Identify underlying intent or goal behind actions	SA-INTENT, OA-INTENT	148	23.0s
Cross-Video (L3, 513)	Sync-Referring	Link corresponding entities across synchronized videos	V1-SA2-V20A, V1-WO2-V2SS, ...	207	94.7s
	Cross-Video Ordering	Determine event order sequence across multiple videos	SA-ORDER-MV, MIX-ORDER-MV, ...	117	110.0s
	POV Identification	Identify who performed what action in which video	SA-POV-ID, OA-POV-ID, ...	189	91.6s

Table 2: Definition of 15 categories of cross-entity referring reasoning questions in GAMEPLAYQA, including the number of questions and average video duration.

they remain prone to hallucination, including fabricating objects, misinterpreting temporal dynamics, and confusing causal relationships (Bai et al., 2024; He et al., 2025; Tu et al., 2025; Seth et al., 2025).

Video Understanding Benchmarks Video understanding benchmarks have evolved from early action recognition datasets toward evaluations emphasizing temporal reasoning, spatial grounding, and long-context comprehension. General video QA benchmarks such as MVBench (Li et al., 2024a), LongVideoBench (Wu et al., 2024), Video-MME (Fu et al., 2025), and MVU-Eval (Peng et al., 2025) assess multimodal models on fine-grained temporal perception and multi-step inference. Domain-specific benchmarks target narrative understanding in movies and TV shows (Tapaswi et al., 2016; Lei et al., 2018). Egocentric benchmarks including Ego4D (Grauman et al., 2022), EgoSchema (Mangalam et al., 2023), ECBench (Dang et al., 2025), and EgoIllusion (Seth et al., 2025) evaluate first-person video understanding and hallucination detection. Embodied QA benchmarks such as OpenEQA (Majumdar et al., 2024) and EmbodiedBench (Yang

et al., 2025) ground reasoning in physical environments. In the video game domain, MarioQA (Mun et al., 2017) pioneered event-centric QA on 2D platformer videos, while recent works explored the feasibility of using MLLMs to detect video game graphics glitches, including GlitchBench (Taesiri et al., 2024), VideoGameQA-Bench (Taesiri et al., 2025), and PhysGame (Cao et al., 2024).

3 The GameplayQA Framework

We collected 3D multiplayer gameplay footage from 9 commercial games spanning diverse genres (see Appendix C for the full game list). Videos were sourced from YouTube, Twitch streams, and existing datasets (Wang et al., 2025). For games requiring synchronized multi-POV footage, we identified groups of streamers who played together in the same match and downloaded their individual recordings, then manually aligned them to construct temporally synchronized multi-video sets.

This section details how we obtain the benchmark from these raw videos: defining a question taxonomy (Section 3.1), annotating via timeline captioning on synchronized multi-POV videos (Section 3.2), generating QA pairs through a com-

binatorial template-based algorithm (Section 3.3), and applying quality assurance procedures (Section 3.4). The final benchmark contains 2.4K QA pairs, generated from 2709 caption true labels and 1586 distractor labels.

3.1 Question Taxonomy

Our question taxonomy (Figure 1) is built upon a **six-primitive label system** that categorizes observable events along two axes: *Agent* (Self, Other, World) and *Temporal Nature* (Action/State for agents, Object/Event for world).

Entity Types. We organize perception in interactive 3D environments around three entity categories (Figure 5): **Self** (the POV agent), **Other** (external entities such as teammates, enemies, or NPCs), and **World** (the shared environment). This Self–Other–World decomposition naturally aligns with multi-agent reinforcement learning frameworks and agent-based modeling paradigms (Sutton and Barto, 2018; Busoniu et al., 2008), where agents must simultaneously track their own state, model other agents’ behaviors, and respond to environmental dynamics (Illustration in Fig. 5). For each entity category, we distinguish between dynamic and static properties: **Self-Action** (SA) captures what the player *does* (shooting, jumping, reloading), while **Self-State** (SS) captures the player’s *condition* (health, ammo, equipped weapon). Similarly, **Other-Action** (OA) and **Other-State** (OS) track other agents. The *World* category is divided into **World-Object** (WO), referring to static or interactive items such as supply crates and vehicles, and **World-Event** (WE), which includes dynamic events like explosions or game notifications. This labeling system enables hallucination analysis of model error rates by entity type (see Sec. 4.2).

Task Categories. We organize questions into 15 task categories across three cognitive levels; question examples, category sizes, and average video durations are summarized in Table 2. **Level 1 (Single Reference)** tests basic perception: recognizing actions, states, objects, and events within a single video segment. These tasks include action recognition (e.g., “What did the player do?”), state recognition (e.g., “What was the player’s health?”), object recognition, event recognition, and static object counting. **Level 2 (Temporal)** introduces temporal reasoning that requires grounding answers to specific time windows. Tasks include cross-entity referring (e.g., “When the player jumped, what was

their health?”), timestamp referring, time localization, absence recognition (identifying what did *not* occur), occurrence counting, temporal ordering, and intent identification. **Level 3 (Cross-Video)** extends reasoning across synchronized multi-POV footage, testing sync-referring (e.g., “When POV1 was reloading, what did POV2 do?”), cross-video ordering, and POV identification. This hierarchy progressively tests from basic perception to complex multi-perspective temporal reasoning. Figure 3 provides typical example questions covering the task categories.

Distractor Taxonomy. A key contribution of GAMEPLAYQA is its structured distractor taxonomy, which enables fine-grained diagnosis of why models hallucinate. We categorize incorrect options by their relationship to the ground truth. **Lexical** distractors are text-based variants of the correct option, generated by changing the subject, using antonyms, or altering object attributes. **Scene** distractors are vision-based options listing plausible events that did not actually occur in the video. **Temporal** distractors refer to events that *did* happen, but outside the queried time window. **Role** distractors swap the agent attribution (e.g., attributing other agents’ actions to the POV player). **Cross-Video** distractors refer to events from other synchronized videos, applicable only to multi-video questions. By analyzing the error rates for each distractor type, we can pinpoint failure modes in temporal grounding, agent attribution, or semantic understanding.

3.2 Multi-Video Timeline Captioning

We employ **dense multi-track timeline captioning** where each of the six entity types (SA, SS, OA, OS, WO, WE) is treated as an independent annotation track (See Figure 7 and Figure 8 for screenshots of labeling interface). Labels within and across tracks can overlap temporally, enabling concurrent event capture (e.g., a player action (SA) occurring while their health state (SS) changes during a world event (WE)). Figure 2 visualizes this process, where the object label “a ladder” is temporally referred to ask a question regarding the player’s action at the same time. For multi-POV videos, we synchronize timelines across perspectives, enabling cross-video temporal alignment.

Decision Density. We operationalize decision density as the temporal frequency of semantic labels such as actions, states, and events that constitute the necessary information stream for an agent’s

planning and reaction loop. Formally, we define the density metric ρ as:

$$\rho = \frac{N_{\text{labels}}}{T_{\text{seconds}}} \quad (1)$$

Across our benchmark, 2,709 true labels span a total of 2,219.41 seconds of annotated footage, yielding $\rho \approx 1.22$ labels/second. Table 10 (Appendix C) shows the per-type breakdown, reflecting the predominance of self-centric observations in first-person gameplay. This high-frequency labeling regime sets GAMEPLAYQA apart from passive video benchmarks and underscores the inherent difficulty of temporal grounding tasks in our experiments.

The annotation process follows a two-stage human-in-the-loop workflow. In the first stage, Gemini-3-Pro generates candidate labels (3,632 predictions) and distractors (1,678 predictions). Four graduate student annotators then verify and refine these candidates: 31.1% of predicted labels were deleted, 42.7% were edited (with 61.9% requiring caption changes and 42.2% requiring temporal boundary adjustments), and 26.2% were accepted without modification. Additionally, 7.6% of the final label set were added entirely by annotators to capture events missed by the model. In the second stage, a separate annotator reviews all labels, making further adjustments to approximately 12% of labels. Detailed annotation protocol and annotator statistics are provided in Appendix E.

3.3 Combinatorial QA Generation

We generate questions through a **combinatorial template-based algorithm** that instantiates question templates by systematically combining verified labels across five orthogonal dimensions: number of videos, context target, entity type, distractor type, and question form, as summarized in Table 2 and Table 7. For each combination, the algorithm selects a ground-truth label as the correct answer and populates the remaining options with distractors drawn from the corresponding distractor pool, enabling fine-grained diagnosis of model failure modes. Complete templates are listed in Appendix F. Optionally, an LLM paraphrasing step is applied to reword the templated questions into more natural phrasing without altering their meaning or answer.

The algorithm initially produces 399,214 candidate QA pairs. Sync-Referring, Cross-Entity Referring, Timestamp Referring, and Ordering types

dominate due to their combinatorial nature, so we strategically downsample to 4K questions to enforce balanced category coverage and avoid long-tail bias. After quality assurance described in Section 3.4, this yields the final 2,365 gold-standard pairs.

3.4 Quality Assurance

Language Prior Filtering. Template-based generation can introduce language priors that allow models to guess answers without visual grounding. To mitigate this, we apply a **blind filtering** procedure: for each generated question, we query Gemini-3-Flash with $k = 3$ trials using only the question text (no video). Questions where the model consistently achieves high accuracy are flagged as potentially biased and removed from the benchmark. This ensures that remaining questions require genuine video understanding rather than exploiting statistical regularities in question phrasing.

Human Evaluation. To validate generation quality, we evenly sampled 120 questions covering all question types for human evaluation. Annotators assessed two criteria: (1) the video contains exactly one correct answer among the options, and (2) the question adheres to the semantics defined by its question code (*e.g.*, an IDENT question truly requires identification). For questions where annotators disagreed, we held discussion meetings to reach consensus; when no agreement could be reached, we resolved through majority voting. During this process, 8% of questions were flagged as faulty due to issues such as excessive similarity between multiple options or misaligned temporal boundaries, which is consistent with the annotation error propagation discussed in our limitations (Section 5).

4 Experiments

We evaluate both open-source and proprietary MLLMs. **Open-source:** Qwen3-VL Series (Bai et al., 2025), Gemma 3 Series (Gemma Team et al., 2025). **Proprietary:** GPT-5 Series (OpenAI, 2025), Claude 4.5 (Anthropic, 2025) (Sonnet, Haiku), Gemini Series (Comanici et al., 2025), and Seed 1.6 (Guo et al., 2025).

Evaluation Setup. We evaluate all models in a zero-shot setting using accuracy as the metric. For video-native models (Gemini, Seed), we input the entire video directly. For frame-based models, we

Model	All	L1 (Single Ref.)					L2 (Temporal)							L3 (Cross-Video)		
		ActRec	StaRec	ObjRec	EvtRec	SOC	X-Ent	TsRef	TimLoc	AbsRec	OccCnt	Order	Intent	SyncRef	X-VOrd	POV-ID
Human	80.5	80.0	80.0	100.0	75.0	100.0	84.2	100.0	76.9	83.3	62.5	77.8	57.1	88.9	77.8	100.0
<i>Proprietary Models</i>																
GPT-5	67.0	79.0	70.7	70.0	68.9	48.3	71.6	70.4	45.9	86.2	62.7	78.3	54.1	72.0	60.7	54.0
GPT-5 Mini	62.7	70.4	67.3	68.6	72.1	34.5	67.6	66.7	47.0	79.0	33.3	72.8	50.0	72.0	43.6	58.7
GPT-5 Nano	49.3	61.7	60.5	70.0	72.1	37.9	57.7	65.4	33.5	64.6	17.3	35.0	41.9	49.8	29.1	42.9
Gemini 2.5 Pro	71.3	69.1	68.0	70.0	80.3	34.5	74.5	77.8	65.1	82.1	38.7	82.8	59.5	81.0	65.0	85.7
Gemini 3 Flash	68.2	71.6	65.3	75.7	68.9	24.1	70.7	80.2	64.4	83.6	32.0	78.9	62.8	76.3	52.1	60.3
Gemini 2.5 Flash	63.7	69.8	59.2	71.4	72.1	31.0	65.0	69.1	60.5	76.9	34.7	72.2	60.8	72.9	50.4	51.3
Claude 4.5 Sonnet	51.3	62.3	49.7	70.0	65.6	48.3	57.9	50.6	34.9	68.2	41.3	42.2	61.5	47.8	30.8	46.0
Claude 4.5 Haiku	41.8	46.9	52.4	60.0	60.7	51.7	41.8	53.1	26.0	53.3	24.0	36.1	46.6	41.1	29.9	38.6
Seed 1.6	61.8	75.9	63.3	77.1	73.8	51.7	70.4	65.4	44.1	78.5	42.7	69.4	60.1	57.0	41.9	47.6
Seed 1.6 Flash	56.5	66.9	56.1	72.1	74.6	50.0	65.5	67.9	30.9	68.8	38.4	41.5	63.1	61.7	48.2	44.2
<i>Open-Source Models</i>																
Qwen3 VL 235B	63.8	71.0	59.9	70.0	80.3	55.2	68.6	76.5	54.4	80.0	50.7	72.8	63.5	66.7	31.6	49.2
Qwen3 VL 30B	60.8	68.5	60.5	74.3	82.0	58.6	65.2	77.8	47.7	79.5	65.3	66.7	56.8	55.1	30.8	47.1
Qwen3 VL 8B	57.8	68.5	56.5	74.3	72.1	62.1	63.6	75.3	46.3	73.3	52.0	57.2	64.9	48.3	27.4	45.5
Gemma 3 27B	48.0	55.6	54.4	58.6	60.7	44.8	57.4	64.2	29.2	66.2	32.0	28.3	50.7	46.4	29.9	46.0
Gemma 3 12B	43.7	53.1	48.3	65.7	59.0	31.0	52.5	54.3	26.7	54.9	9.3	27.2	50.0	50.2	24.8	39.7
Gemma 3 4B	42.9	46.9	42.9	64.3	63.9	24.1	49.6	54.3	26.0	57.4	9.3	27.2	46.6	58.5	23.9	37.6
<i>All Models</i>	56.9	64.8	58.4	69.5	70.4	43.0	62.5	66.8	42.7	72.0	36.5	55.5	55.8	59.8	38.8	49.6

Table 3: Model performance across task categories. **Gold** = best, **Silver** = second best. L1: ActRec (Action Recognition), StaRec (State Recognition), ObjRec (Object Recognition), EvtRec (Event Recognition), SOC (Static Object Count). L2: X-Ent (Cross-Entity Referring), TsRef (Timestamp Referring), TimLoc (Time Localization), AbsRec (Absence Recognition), OccCnt (Occurrence Count), Order (Ordering), Intent (Intent Identification). L3: SyncRef (Sync-Referring), X-VOrd (Cross-Video Ordering), POV-ID (POV Identification).

sample frames at 1 FPS up to 32 frames; for videos longer than 32 seconds, we uniformly sample 32 frames across the duration. Videos are resized such that the longer side is 720p while preserving aspect ratio. Although models are instructed to output a single letter, they sometimes produce full sentences or explanations; we use GPT-5-mini as an LLM judge to extract the selected option. Detailed inference settings are in Appendix B; evaluation prompt templates in Appendix D.

4.1 Main Results

Table 3 summarizes model performance across all task categories. Among all models evaluated, Gemini 2.5 Pro attains the highest overall accuracy (71.3%), followed by Gemini 3 Flash (68.2%) and GPT-5 (67.0%), yet a substantial gap to human performance (80.5%) persists. We highlight two key findings below.

Consistent degradation across cognitive levels. Averaged across all models, accuracy drops steadily from L1 Single-Reference (61.2%) to L2 Temporal (56.0%) to L3 Cross-Video (49.4%). This trend validates that the three-level hierarchy of GAMEPLAYQA successfully stratifies task difficulty, with temporal grounding and multi-POV reasoning remaining substantially more challenging than basic visual perception.

Counting and Cross-Video Ordering are the hardest tasks. Two tasks emerge as clear bottlenecks. Occurrence Count (OccCnt) averages only 36.5% across models, making it the hardest L2 task. This suggests that tracking event recurrences over time, which demands sustained temporal attention across frames, remains beyond the reach of current models. Cross-Video Ordering (X-VOrd) averages 38.8%, the lowest among L3 tasks, with several models dropping to around 30%, indicating severe difficulty in aligning temporal events across perspectives. Together, these results suggest that precise temporal tracking, whether within a single video or across multiple perspectives, remains a fundamental weakness of current video-language architectures.

4.2 Error Source Analysis

We conduct a fine-grained error analysis to identify systematic failure modes by entity category. Table 4 reveals that **World-Object (WO)** recognition is the easiest category (62.0% aggregate), while recognizing **Other Agents** proves substantially harder, for example Other-Action (OA) at 54.0% and Other-State (OS) at 55.4% represent an 8-point gap compared to world objects. This suggests MLLMs struggle with other agent attribution in multi-agent scenes.

We further plot error rates along four criteria: dis-

Model	All	SA	SS	OA	OS	WO	WE
<i>Proprietary Models</i>							
Gemini 2.5 Pro	71.3	75.2	73.3	65.6	72.0	71.7	74.0
Gemini 3 Flash	68.2	70.5	73.0	65.6	64.8	70.5	68.4
GPT-5	67.0	66.0	73.5	60.8	67.0	71.6	69.1
Gemini 2.5 Flash	63.7	65.1	66.6	60.8	60.1	67.6	64.9
GPT-5 Mini	62.7	64.4	68.6	57.6	60.4	66.8	63.8
Seed 1.6	61.8	60.3	65.8	57.6	62.3	67.0	65.6
Seed 1.6 Flash	56.5	56.9	59.3	56.7	49.6	63.0	63.5
Claude 4.5 Sonnet	51.3	51.6	53.0	50.9	46.9	57.8	53.6
GPT-5 Nano	49.3	47.1	59.9	47.7	50.3	55.9	51.4
Claude 4.5 Haiku	41.8	41.0	39.3	38.1	44.0	44.1	45.7
<i>Open-Source Models</i>							
Qwen3 VL 235B	63.8	63.6	67.4	59.2	58.8	70.8	70.2
Qwen3 VL 30B	60.8	57.6	62.7	57.1	60.1	66.2	67.8
Qwen3 VL 8B	57.8	53.7	59.9	56.8	56.3	65.9	63.4
Gemma 3 27B	48.0	48.4	53.5	48.3	47.8	54.6	49.7
Gemma 3 12B	43.7	45.2	51.4	42.1	44.7	47.3	46.4
Gemma 3 4B	42.9	43.4	49.6	43.2	40.3	52.2	47.7
<i>All Models</i>	56.8	56.5	61.0	54.0	55.4	62.0	60.2

Table 4: Model performance by entity category. **Gold** = best, **Silver** = second best. SA: Self-Action, SS: Self-State, OA: Other-Action, OS: Other-State, WO: World-Object, WE: World-Event.

tractor type in EXIST questions (True/False), game name, video length, and number of synchronized videos. Figure 4 reveals three key insights. First, models are primarily confused by cross-video and temporal distractors, while scene distractors are the easiest, indicating that models handle static visual input better than temporal and cross-video reasoning. Second, game pace strongly predicts difficulty: competitive shooters with rapid state transitions (Counter-Strike, Battlefield, Apex Legends) rank as the top three hardest games compared to slower exploration titles, validating that *decision-dense environments* pose fundamentally harder challenges. Third, both temporal extent and multi-POV complexity compound errors, as longer clips and additional synchronized perspectives each generally degrade performance monotonically.

4.3 Language Prior and Temporal Ablation

To disentangle the contributions of visual grounding and temporal reasoning, we conduct ablation studies on GPT-5-mini under three degraded input conditions: (1) **No Video**, where only the question text is provided; (2) **Random Frame**, where a single randomly selected frame replaces the full video; and (3) **Shuffled Frames**, where the original frames are presented in random order. Results are shown in Table 5.

The No Video condition drops accuracy the largest amount (33.3%), confirming that GAMEPLAYQA requires genuine visual grounding and

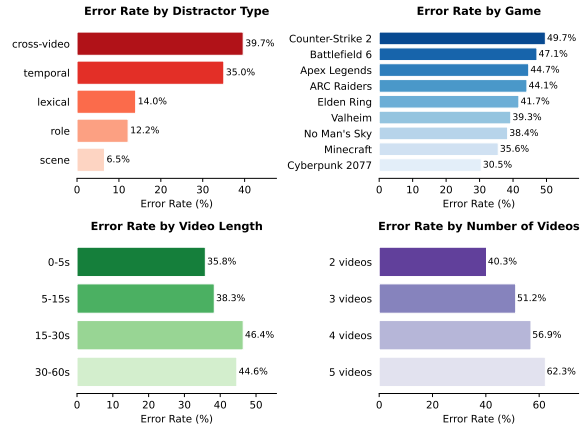


Figure 4: Error rate analysis across four dimensions. **Top-left:** Cross-video and temporal distractors cause the most errors. **Top-right:** Fast-paced shooters (CS2, Battlefield) are hardest. **Bottom-left:** Error increases with video length. **Bottom-right:** Error scales with number of synchronized videos.

Condition	All	L1	L2	L3
Baseline (Full)	62.7	67.2	61.9	60.6
No Video	29.4	36.0	29.1	24.2
Random Frame	41.7	52.9	40.9	33.7
Shuffled Frames	54.8	63.1	52.6	53.4

Table 5: Ablation study on GPT-5-mini with degraded visual inputs. Performance drops indicate the contribution of video content and temporal ordering.

cannot be solved by language priors alone. The Random Frame condition recovers only partial performance (+12.3% over No Video), indicating that static visual content provides useful context but cannot substitute for temporal dynamics. Shuffled Frames achieves near-baseline L1 performance (63.1% vs. 67.2%) but degrades substantially on L2 (52.6% vs. 61.9%) and L3 (53.4% vs. 60.6%), showing that temporal ordering is critical for reasoning tasks but less so for basic perception.

4.4 Cross-Domain Generalization

To validate that our framework generalizes beyond gameplay to broader single-agent and multi-agent ego-centric settings, we conducted a small-scale transfer experiment by applying the identical benchmarking pipeline to two real-world domains: (1) dashcam collision videos from the Nexar dataset (Moura et al., 2025), and (2) synchronized ego-centric videos of humans collaboratively assembling Lego from the Ego-Humans benchmark (Khirodkar et al., 2023). The only pipeline adjustment required was renaming the de-

Model	All	L1 (Single Ref.)				L2 (Temporal)						L3 (Cross-Video)			
		ActRec	StaRec	ObjRec	EvtRec	X-Ent	TsRef	TimLoc	AbsRec	OccCnt	Order	Intent	SyncRef	X-VOrd	POV-ID
		Gemini 2.5 Pro	66.2	61.1	86.7	83.3	100.0	62.0	80.0	52.6	72.2	40.0	50.0	66.7	76.7
Gemini 2.5 Flash	62.0	72.2	73.3	66.7	80.0	64.0	80.0	36.8	61.1	40.0	83.3	55.6	70.0	50.0	38.9
GPT-5 Mini	61.0	55.6	66.7	100.0	80.0	66.0	70.0	63.2	77.8	20.0	66.7	66.7	53.3	50.0	27.8
Qwen3 VL 235B	59.2	50.0	73.3	100.0	100.0	66.0	80.0	36.8	72.2	20.0	50.0	66.7	50.0	25.0	44.4

Table 6: Cross-domain generalization results on real-world ego-centric videos (autonomous driving + multi-human collaboration, 213 questions). **Gold** = best, **Silver** = second best. Column headers follow the same conventions as Table 3.

fault actor from “player” to domain-appropriate labels such as “person” or “driver”.

Across 4 videos (~113 seconds total), our automated pipeline generated 5,463 initial questions; following the same downsampling and quality-assurance protocol as the main benchmark, we produced a test set of **213 questions** at a label density of $\rho = 0.50$ labels/second, lower than gameplay ($\rho = 1.22$), reflecting the slower decision pace of real-world activities.

Table 6 shows that performance trends mirror those of the main benchmark: Gemini 2.5 Pro leads overall (66.2%), models degrade progressively from L1 to L3, and Occurrence Count and Cross-Video Ordering remain the hardest tasks. The lower label density confirms that real-world videos progress at a slower decision pace than gameplay, yet the relative difficulty ordering across models and task categories is preserved. These results demonstrate that our benchmarking framework generalizes to real-world spatiotemporal tasks with only minimal domain-specific adjustments.

5 Conclusion

We presented GAMEPLAYQA, an end-to-end benchmarking framework that uses densely annotated, synchronized multi-POV gameplay videos to evaluate agentic perception in decision-dense 3D environments. Built on a Self–Other–World entity decomposition and a three-level cognitive hierarchy, the framework refines 2.4K diagnostic QA pairs whose structured distractors pinpoint where models hallucinate. Evaluation of 16 frontier MLLMs reveals steady performance degradation from basic perception to temporal reasoning to cross-video understanding, with models particularly failing on other-agent attribution, temporal grounding, and fast-paced decision-dense scenarios. Cross-domain experiments on autonomous driving and ego-centric human collaboration confirm that the pipeline generalizes with minimal

adaptation, preserving difficulty and model rankings across domains. We hope GAMEPLAYQA drives progress toward MLLMs capable of reliable perception and reasoning in dynamic, multi-agent worlds.

Limitations

While GAMEPLAYQA includes intent identification as a proxy for understanding goal-directed behavior, our framework does not cover decision reasoning questions such as “*What is the best action to take at this moment?*”. Answering such questions would require estimating expected rewards or action values from raw video observations, which is a capability that remains an open research challenge, as it demands not only perception but also learning implicit reward structures and world dynamics from uncurated, in-the-wild footage. Additionally, intent identification is inherently more subjective than recognizing physical actions or states, which occasionally results in multiple defensible answers; during human evaluation, approximately 8% of questions were flagged as having ambiguous ground-truth labels (Section 3.4). Nevertheless, because anticipating intent is a critical capability for planning agents, we believe this category provides vital diagnostic signal and should be preserved.

Annotation Cost and Error Propagation. The dense labeling process underlying GAMEPLAYQA is extremely labor-intensive and susceptible to human error. Annotators must track over 100 labels and distractors per video, repeatedly watching the same footage while switching cognitive focus across different entity types and temporal windows. On average, labeling a 30-second video clip requires 25–35 minutes. More critically, the combinatorial QA generation algorithm reuses labels across multiple questions, meaning that a single labeling error, whether in timestamp boundaries, entity type, or description content, can propagate to multiple erroneous questions. However, because the QA

generation algorithm is deterministic and template-based, perfectly accurate annotations would yield perfectly correct questions; i.e. the source of benchmark noise is annotation error, not algorithmic error. While our quality assurance procedures mitigate this risk, some annotation noise inevitably remains in the benchmark.

Acknowledgments

The authors acknowledge the use of Large Language Models for assistance with proofreading and grammar checking. All content was reviewed, edited, and approved by the human authors, who take full responsibility for the final manuscript. The project or effort depicted was or is sponsored by the U.S. Army Combat Capabilities Development Command – Soldier Centers under contract number W912CG-24-D-0001. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

References

- Anthropic. 2025. [System card: claude sonnet 4.5](#). Technical report, Anthropic.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, and 29 others. 2023. [Qwen technical report](#). *Preprint*, arXiv:2309.16609.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025. [Qwen3-vl technical report](#). *Preprint*, arXiv:2511.21631.
- Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and Mike Zheng Shou. 2024. Hallucination of multimodal large language models: A survey. *arXiv preprint arXiv:2404.18930*.
- Lucian Busoniu, Robert Babuska, and Bart De Schutter. 2008. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2):156–172.
- Meng Cao, Haoran Tang, Haoze Zhao, Hangyu Guo, Jiaheng Liu, Ge Zhang, Ruyang Liu, Qiang Sun, Ian Reid, and Xiaodan Liang. 2024. Physgame: Uncovering physical commonsense violations in gameplay videos. *arXiv preprint arXiv:2412.01800*.
- Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, and Lidong Bing. 2024. [Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms](#). *Preprint*, arXiv:2406.07476.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szpektor, Nan-Jiang Jiang, and 3416 others. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#). *Preprint*, arXiv:2507.06261.
- Ronghao Dang, Yuqian Yuan, Wenqi Zhang, Yifei Xin, Boqiang Zhang, Long Li, Liuyi Wang, Qinyang Zeng, Xin Li, and Lidong Bing. 2025. Ecbench: Can multimodal foundation models understand the egocentric world? a holistic embodied cognition benchmark. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 24593–24602.
- Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, and 1 others. 2025. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 24108–24118.
- Gemini Robotics Team, Abbas Abdolmaleki, Saminda Abeyruwan, Joshua Ainslie, Jean-Baptiste Alayrac, Montserrat Gonzalez Arenas, Ashwin Balakrishna, Nathan Batchelor, Alex Bewley, Jeff Bingham, Michael Bloesch, Konstantinos Bousmalis, Philemon Brakel, Anthony Brohan, Thomas Buschmann, Arunkumar Byravan, Serkan Cabi, Ken Caluwaerts, Federico Casarini, and 153 others. 2025. [Gemini robotics 1.5: Pushing the frontier of generalist robots with advanced embodied reasoning, thinking, and motion transfer](#). *Preprint*, arXiv:2510.03342.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. [Gemma 3 technical report](#). *Preprint*, arXiv:2503.19786.
- Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, Miguel Martin, Tushar Nagarajan, Ilija Radosavovic, Santhosh Kumar Ramakrishnan, Fiona Ryan, Jayant Sharma, Michael Wray, Mengmeng Xu, Eric Zhongcong Xu, and 66 others. 2022. Ego4d: Around the world in 3,000 hours of egocentric video.

- In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18995–19012.
- Dong Guo, Faming Wu, Feida Zhu, Fuxing Leng, Guang Shi, Haobin Chen, Haoqi Fan, Jian Wang, Jianyu Jiang, Jiawei Wang, Jingji Chen, Jingjia Huang, Kang Lei, Liping Yuan, Lishu Luo, Pengfei Liu, Qinghao Ye, Rui Qian, Shen Yan, and 178 others. 2025. [Seed1.5-vl technical report](#). *Preprint*, arXiv:2505.07062.
- Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Yong Dai, Hongming Zhang, Zhenzhong Lan, and Dong Yu. 2024. Webvoyager: Building an end-to-end web agent with large multimodal models. *arXiv preprint arXiv:2401.13919*.
- Yixiao He, Haifeng Sun, Pengfei Ren, Jingyu Wang, Huazheng Wang, Qi Qi, Zirui Zhuang, and Jing Wang. 2025. Evaluating and mitigating object hallucination in large vision-language models: Can they still see removed objects? In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6841–6858.
- Pablo Hernandez-Leal, Bilal Kartal, and Matthew E Taylor. 2019. A survey and critique of multiagent deep reinforcement learning. *Autonomous Agents and Multi-Agent Systems*, 33(6):750–797.
- Rawal Khirodkar, Aayush Bansal, Lingni Ma, Richard Newcombe, Minh Vo, and Kris Kitani. 2023. Ego-humans: An ego-centric 3d multi-human benchmark. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19807–19819.
- Jie Lei, Licheng Yu, Mohit Bansal, and Tamara Berg. 2018. Tvqa: Localized, compositional video question answering. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 1369–1379.
- Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, Limin Wang, and Yu Qiao. 2024a. [Mvbench: A comprehensive multimodal video understanding benchmark](#). *Preprint*, arXiv:2311.17005.
- Yanwei Li, Chengyao Wang, and Jiaya Jia. 2024b. Llama-vid: An image is worth 2 tokens in large language models. In *European Conference on Computer Vision*, pages 323–340. Springer.
- Qian Long, Zhi Li, Ran Gong, Ying Nian Wu, Demetri Terzopoulos, and Xiaofeng Gao. 2024. Teamcraft: A benchmark for multi-modal multi-agent systems in minecraft. *arXiv preprint arXiv:2412.05255*.
- Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mcvay, Oleksandr Maksymets, Sergio Arnaud, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent Berges, Shiqi Zhang, Pulkit Agrawal, Yonatan Bisk, Dhruv Batra, and 5 others. 2024. Openeqa: Embodied question answering in the era of foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16488–16498.
- Karttikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. 2023. Egoschema: A diagnostic benchmark for very long-form video language understanding. *Advances in Neural Information Processing Systems*, 36:46212–46244.
- Daniel Moura, Shizhan Zhu, and Orly Zviti. 2025. Nexar dashcam collision prediction dataset and challenge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2583–2591.
- Jonghwan Mun, Paul Hongsuck Seo, Ilchae Jung, and Bohyung Han. 2017. [MarioQA: Answering Questions by Watching Gameplay Videos](#). In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2886–2894, Los Alamitos, CA, USA. IEEE Computer Society.
- OpenAI. 2025. [Gpt-5 system card](#). Technical report, OpenAI.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Tianhao Peng, Haochen Wang, Yuanxing Zhang, Zekun Wang, Zili Wang, Gavin Chang, Jian Yang, Shihao Li, Yanghai Wang, Xintao Wang, Houyi Li, Wei Ji, Pengfei Wan, Steven Huang, Zhaoxiang Zhang, and Jiaheng Liu. 2025. [Mvu-eval: Towards multi-video understanding evaluation for multimodal llms](#). *Preprint*, arXiv:2511.07250.
- Georgy Savva, Oscar Michel, Daohan Lu, Suppakit Waiwitlikhit, Timothy Meehan, Dhairya Mishra, Srivats Poddar, Jack Lu, and Saining Xie. 2026. Solaris: Building a multiplayer video world model in minecraft. *arXiv preprint arXiv:2602.22208*.
- Ashish Seth, Utkarsh Tyagi, Ramaneswaran Selvakumar, Nishit Anand, Sonal Kumar, Sreyan Ghosh, Ramani Duraiswami, Chirag Agarwal, and Dinesh Manocha. 2025. Egoillusion: Benchmarking hallucinations in egocentric video understanding. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 28449–28468.
- SIMA team, Adrian Bolton, Alexander Lerchner, Alexandra Cordell, Alexandre Moufarek, Andrew Bolt, Andrew Lampinen, Anna Mitenkova, Arne Olav Hallingstad, Bojan Vujatovic, Bonnie Li, Cong Lu, Daan Wierstra, Daniel P. Sawyer, Daniel

- Slater, David Reichert, Davide Vercelli, Demis Hassabis, Drew A. Hudson, and 47 others. 2025. [Sima 2: A generalist embodied agent for virtual worlds](#). *Preprint*, arXiv:2512.04797.
- SIMA Team, Maria Abi Raad, Arun Ahuja, Catarina Barros, Frederic Besse, Andrew Bolt, Adrian Bolton, Bethanie Brownfield, Gavin Buttmore, Max Cant, Sarah Chakera, Stephanie C. Y. Chan, Jeff Clune, Adrian Collister, Vikki Copeman, Alex Culum, Ishita Dasgupta, Dario de Cesare, Julia Di Trapani, and 75 others. 2024. [Scaling intractable agents across many simulated worlds](#). *Preprint*, arXiv:2404.10179.
- Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*, second edition. MIT Press.
- Mohammad Reza Taesiri, Tianjun Feng, Cor-Paul Bezeemer, and Anh Nguyen. 2024. Glitchbench: Can large multimodal models detect video game glitches? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22444–22455.
- Mohammad Reza Taesiri, Abhijay Ghildyal, Saman Zadtootaghaj, Nabajeet Barman, and Cor-Paul Bezeemer. 2025. [Videogameqa-bench: Evaluating vision-language models for video game quality assurance](#). In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2016. Movieqa: Understanding stories in movies through question-answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4631–4640.
- Yahan Tu, Rui Hu, and Jitao Sang. 2025. Ode: Open-set evaluation of hallucinations in multimodal large language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19836–19845.
- Yunzhe Wang, Soham Hans, and Volkan Ustun. 2025. X-ego: Acquiring team-level tactical situational awareness via cross-egocentric contrastive video representation learning. *arXiv preprint arXiv:2510.19150*.
- Haoning Wu, Dongxu Li, Bei Chen, and Junnan Li. 2024. Longvideobench: A benchmark for long-context interleaved video-language understanding. *Advances in Neural Information Processing Systems*, 37:28828–28857.
- Rui Yang, Hanyang Chen, Junyu Zhang, Mark Zhao, Cheng Qian, Kangrui Wang, Qineng Wang, Teja Venkat Koripella, Marziyeh Movahedi, Manling Li, Heng Ji, Huan Zhang, and Tong Zhang. 2025. [Embodiedbench: Comprehensive benchmarking multi-modal large language models for vision-driven embodied agents](#). *Preprint*, arXiv:2502.09560.
- Yuguang Yue, Irakli Salia, Samuel Hunt, Chris Green, Wenzhe Shi, and Jonathan J Hunt. 2026. Scaling behavior cloning improves causal reasoning: An open model for real-time video game playing. *arXiv preprint arXiv:2601.04575*.
- Chaoyun Zhang, Liqun Li, He Huang, Chiming Ni, Bo Qiao, Si Qin, Yu Kang, Minghua Ma, Qingwei Lin, Saravan Rajmohan, and Dongmei Zhang. 2026. [Ufo3: Weaving the digital agent galaxy](#). *Preprint*, arXiv:2511.11332.
- Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, Quan Vuong, Vincent Vanhoucke, Huong Tran, Radu Soricut, Anikait Singh, Jaspiar Singh, Pierre Sermanet, Pannag R. Sanketi, Grecia Salazar, and 35 others. 2023. [Rt-2: Vision-language-action models transfer web knowledge to robotic control](#). In *Proceedings of The 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pages 2165–2183. PMLR.

A GameplayQA Question Taxonomy



Figure 5: Illustration of the **Self-Other-World** framework. A first-person player (Self) perceives a teammate (Other) issuing a warning, set against the surrounding game environment (World). These three perspectives define the entity types used in our question taxonomy.

Table 7 summarizes the GameplayQA question taxonomy. Each question is defined by five orthogonal dimensions. For each category, we provide a short description and representative example questions. Figure 5 illustrates the Self-Other-World concept.

B Model Details and Inference Settings

This section provides details about the inference settings and providers used to run the benchmark. Table 8 summarizes the inference configurations for all tested models.

Video Input Strategy. For **video-native models** (Gemini-2.5-Pro, Gemini-3-Flash, Gemini-2.5-Flash), we input the entire video directly without frame sampling. For **non-video-native models** (GPT, Claude, Qwen, Gemma), we sample frames at 1 FPS up to a default of 32 frames; for videos longer than 32 seconds, we uniformly sample 32 frames across the duration. All videos are resized such that the longer side is 720p while preserving aspect ratio.

Special Cases. Qwen 30B and 235B: Due to provider (Fireworks) limitations, we cap the total number of frames at 30 for these models. For multi-video questions, the 30 frames are evenly split across all input videos (e.g., 10 frames per video for 3 synchronized videos).

Seed 1.6 and Seed 1.6 Flash: Although these models support video-native input, we experienced high instability when calling the API with raw video. We therefore resort to 32-frame sampling for these models to ensure consistent evaluation.

Gemini 3-Pro: We did not benchmark Gemini-3-Pro due to strict rate limits (250 API calls per day) at the time this paper was written, which made full benchmark evaluation infeasible.

Reasoning Effort Settings. For models that support configurable reasoning modes, we use the default reasoning effort settings provided by each API. Table 9 summarizes the default reasoning modes for each model family.

C Dataset Statistics

C.1 Games Selection and Data Source

GameplayQA includes footage from 9 commercially released games spanning diverse genres:

- **Single-POV Games:** Minecraft, Apex Legends, No Man’s Sky, Elden Ring, Cyberpunk 2077, Valheim
- **Multi-POV Synchronized Games:** Counter-Strike 2 (Wang et al., 2025), Battlefield 6, Arc Raiders

For multi-POV games, synchronized footage was obtained either from existing datasets or by identifying groups of Twitch streamers who played together in the same match and manually aligning their individual recordings.

C.2 Label Distribution

Table 10 reports the per-type breakdown of the 2,709 true labels annotated across 2,219.41 seconds of footage, corresponding to a decision density of $\rho \approx 1.22$ labels/second.

Label Type	Count	Share
Self-Action (SA)	658	24.3%
Self-State (SS)	729	26.9%
Other-Action (OA)	160	5.9%
Other-State (OS)	190	7.0%
World-Event (WE)	417	15.4%
World-Object (WO)	555	20.5%
Total	2,709	100%

Table 10: Per-type label distribution ($\rho \approx 1.22$ labels/s over 2,219.41 s of annotated footage).

Figure 6 provides a broader view of the dataset regarding the distribution of question opening phrases, the question codes distribution by count and task type, and word distribution by entity type.

D Prompt Templates

Single-Video Question Answering

Watch the video carefully and answer the following multiple choice question:

<frame_1> <frame_2> ... <frame_32>

Q: <question>

<options>

Please select the correct answer from the options. Answer with the letter directly.

Your answer:

Multi-Video Question Answering

Watch the video carefully and answer the following multiple choice question:

The following are 32 frames of the Video 1:

<frame_1> <frame_2> ... <frame_32>

The following are 32 frames of the Video 2:

<frame_1> <frame_2> ... <frame_32>

...

Q: <question>

<options>

Please select the correct answer from the options. Answer with the letter directly.

Your answer:

LLM as a Judge (Extract Selected Option)

You judge which option a model selected for a multiple choice question.

The question was:

<question>

Available options are: <options>

The model's response was:

<model_output>

Your task is to determine which option the model selected. Look for:

- Explicit mention of a letter (e.g., "A", "B", "C", "D")
- The model stating or implying a specific choice
- The response content matching one of the available options

If the model clearly selected one of the options, return **the corresponding letter**.

If the model's response is empty, an error, unclear, or does not make a definitive choice, return **"X"**.

Language Prior Filtering

You are answering multiple choice questions about video game footage.

You have NOT seen the video. Based only on the question and options provided, select the most likely answer.

You must respond with ONLY a single letter (A, B, C, or D).

E Annotation Protocol

E.1 Annotator Demographics and Expertise

Our annotation team consisted of 5 graduate students (ages 21–31; 3 male, 2 female), assigned to 4 labeler and 2 evaluator roles, with one participant serving in both capacities to ensure cross-stage consistency. Annotators were graduate student co-authors recruited internally for this study and were not financially compensated. All annotators were informed of the purpose of the data collection and consented to their annotations being used for research and potential public release.

General Gaming Experience. To confirm that annotators were qualified to interpret complex game states, we surveyed their overall gaming habits and game-specific familiarity before the annotation process.

- **Video game play frequency in the past 100 days:** 60% (3 participants) play regularly (3–5 times/week); 40% (2 participants) play occasionally (1–2 times/week).
- **Years of video game experience:** 60% (3 participants) have 8+ years; 20% (1 participant) 3–8 years; 20% (1 participant) 1–3 years.

Game-Specific Familiarity. Table 11 reports each annotator's self-reported familiarity with the benchmark titles.

Game	Expert (>300h)	Regular (>30h)	Casual (>5h)	Low (>0h)	None
Counter-Strike	1	1	2	1	0
Minecraft	1	1	3	0	0
Apex Legends	0	2	0	3	0
Battlefield 6	0	2	0	3	0
Arc Raiders	0	1	0	4	0
Cyberpunk 2077	0	1	1	3	0
Elden Ring	0	1	1	2	1
Valheim	0	0	1	3	1

Table 11: Annotator familiarity with each game. Values indicate the number of annotators (out of 5) at each familiarity level. Hours denote minimum play time.

E.2 Annotation Interface

We developed a custom web-based annotation tool supporting both single-video and multi-video labeling. The annotation starts with Gemini-3-Pro-generated timeline captions on the video about the six target entity types (SA, SS, OA, OS, WO, WE) and distractor candidates for lexical semantic distractors and scene distractors. Human annotators then verify each individual caption and distractor regarding type, content, and timeline. The annotation tool is publicly available:

- **Source Code:** <https://github.com/wangyz1999/sync-video-label>
- **Demo Video:** <https://www.youtube.com/watch?v=PKedELJ4XT0>
- **Live Demo:** <https://sync-video-label.vercel.app/>

Figure 7 shows an example of the interface in the single-video setting in the game *Valheim*, where the player is combating a Boar. Figure 8 shows an example of the interface in the multi-video setting in the game *Arc Raiders*, where an explosion in the sky is synchronously captured by all three videos.

E.3 Annotation Instructions

The annotation process follows a structured workflow designed to ensure temporal precision and semantic accuracy. Annotators begin by importing synchronized video instances (single or multi-POV) into the annotation interface. The process proceeds through three sequential phases: *label generation*, *verification*, and *question preview*.

Label Generation and Verification. Annotators initiate automated caption generation using Gemini-3-Pro, which produces candidate true labels, lexical distractors, and scene distractors for each video segment. For each generated true label, annotators verify three criteria: (1) the described event actually occurred, (2) the temporal boundaries $[t_{start}, t_{end}]$ accurately delimit the event, and (3) the assigned entity type (SA, SS, OA, OS, WO, WE) correctly categorizes the label. Temporal boundaries are adjusted when necessary to ensure precise alignment with observable events. For lexical distractors, annotators confirm that the described events do *not* occur during the specified video segment. For scene distractors, verification ensures non-occurrence across the entire video duration. Incorrectly generated labels are not immedi-

ately discarded; instead, we first consider whether they can be repurposed as distractors.

Overall, 31.1% of Gemini-3-Pro-predicted labels were deleted as incorrect or irreparable, while 42.7% were edited to meet quality standards. Of these edited labels, 61.9% required caption text corrections (e.g., fixing entity names, action descriptions, or semantic precision) and 42.2% required temporal boundary adjustments to better align $[t_{start}, t_{end}]$ with the observable event. The remaining 26.2% of predicted labels were accepted without modification. Additionally, 7.6% of the final label set were added entirely by annotators to capture events missed by the model. In the second stage, a separate evaluator reviewed all labels and made further adjustments to approximately 12% of labels, primarily to enforce cross-video consistency and resolve edge cases.

Count Labeling. Count questions require special handling: for action and event types (SA, OA, WE), annotators mark multiple temporally distinct segments sharing the same caption name. For object types (WO), annotators specify the quantity directly in the label metadata.

Quality Control. Annotators systematically review labels by filtering by entity type, verifying each category independently. Ambiguous labels, those where truth value cannot be reliably determined, are removed to maintain benchmark integrity. The interface supports efficient navigation through keyboard shortcuts and contextual menus, enabling rapid verification of temporal alignment and semantic correctness. Once all labels and distractors are verified, annotators proceed to the question preview phase, where generated questions are validated against the verified label set.

F Question Templates

This section provides the complete list of question templates for all three levels used for question generation. Each template is identified by a code combining the entity type and question form. Placeholders: `{other}` refers to other players (teammate, enemy, NPC); `{caption}` refers to specific action/state/object/event descriptions; `{refCaption}` refers to the anchor entity description; `{timestamp}` refers to a formatted time range (e.g., `[00:01 to 00:12]`).

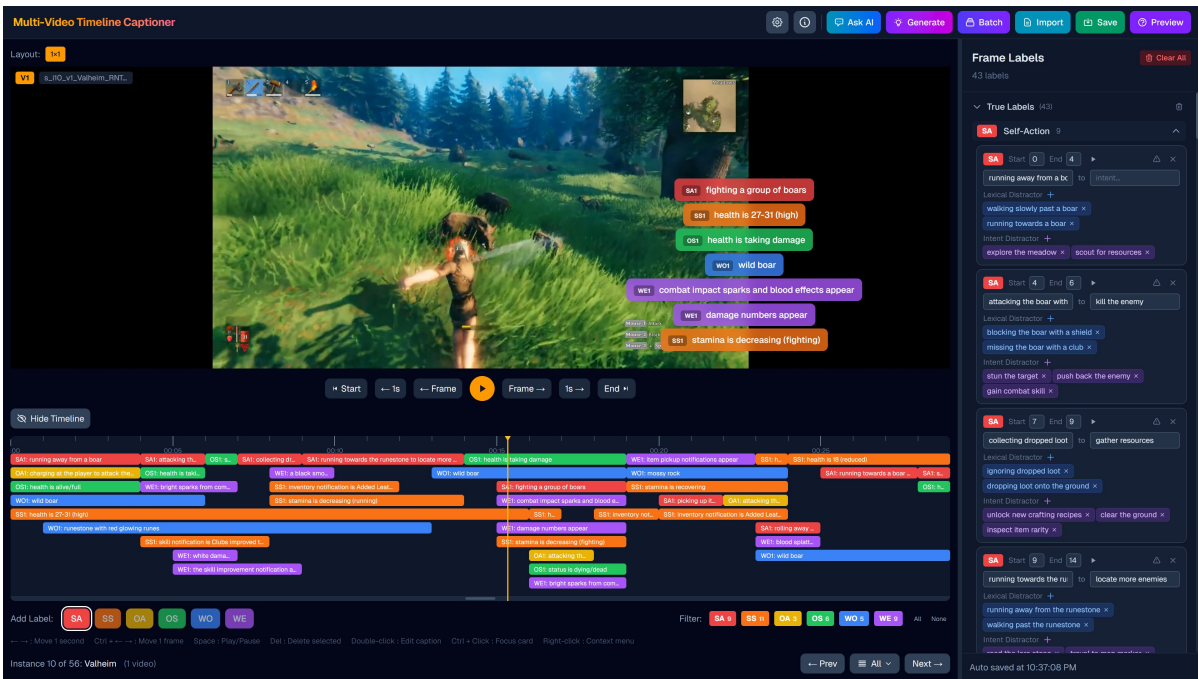


Figure 7: Screenshot of the annotation interface in the **single-video** setting, shown for the game *Valheim*. Annotators label actions, states, events, and entities on a synchronized timeline.

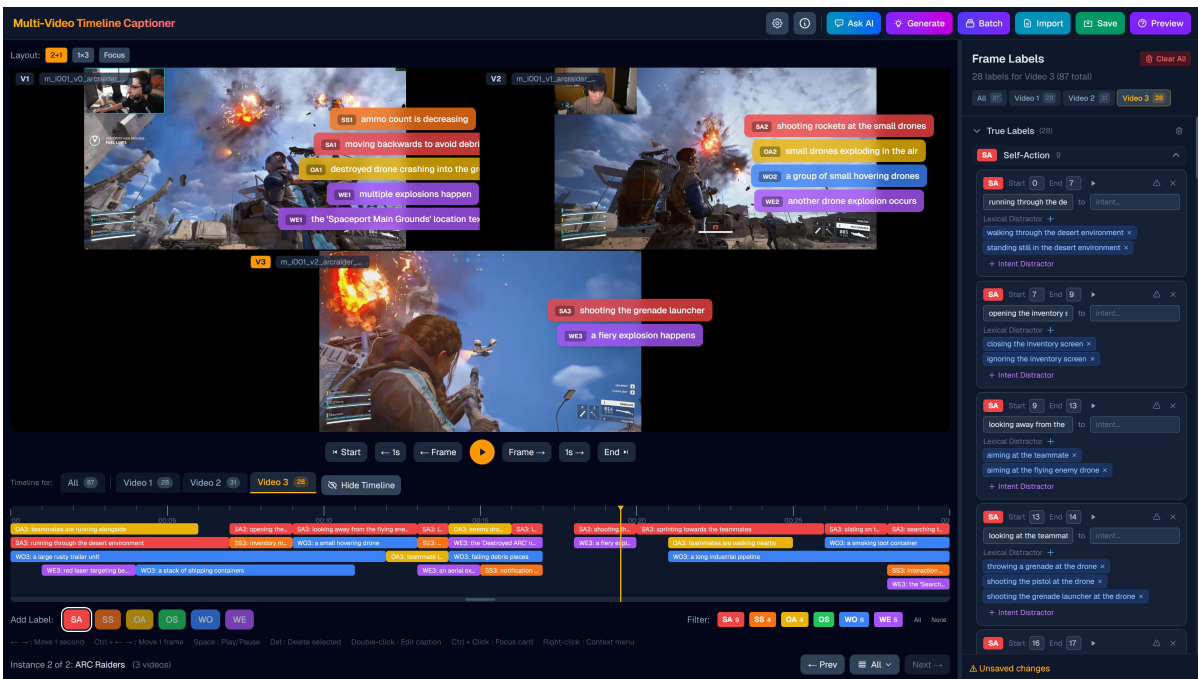


Figure 8: Screenshot of the annotation interface in the **multi-video** setting, shown for the game *Arc Raiders*. Three synchronized video perspectives are displayed in parallel, with aligned timelines enabling annotators to capture cross-video events and temporal relations. The same explosion in the sky is captured by all three videos.

Dimension	Category	Description	Example Question
<i>Number of Videos</i>	Single Video	Inputs a single video	What action did the player perform?
	Multi-Video	Inputs multiple videos	When POV1 player was reloading, which action did POV2 player perform?
<i>Context Target</i>	Summative	(Single-Video Only) Requires aggregation over a temporal segment.	Which of the following best summarizes the player’s actions during the video?
	Timestamp Referring	Refers to a specific moment or interval	At [02:45 - 02:52], what action did the player perform?
	Target Entity Referring	Refers to the temporal moment of the 6 entity types.	When the player was reloading, what action did the teammate perform?
	Cross-Video Referring	(Multi-Video Only) Explicitly references a video index.	Which POV player reload their weapon?
<i>Entity Type</i>	Self-Action	Action performed by the POV player <i>e.g.</i> , shooting, reloading, using item, etc	Which action did the POV player performed?
	Self-State	State or status of the POV player <i>e.g.</i> , health, inventory, equipped weapon, etc	What was the player’s health status at that time?
	Other-Action	Action performed by another player <i>e.g.</i> , teammate, enemy, NPC, etc	What action did the enemy perform during the fight?
	Other-State	State of another player	Was the teammate downed during the encounter?
	World-Object	Environment objects or landmarks <i>e.g.</i> , tree, supply crate, building, cars, etc	How many supply crates are visible in the scene?
	World-Event	Environmental or system-level events or notifications <i>e.g.</i> , explosion, enemy downed, achievement notification, etc	Did an explosion occur during this segment?
<i>Distractor Type</i>	Lexical	Textually similar but incorrect descriptions	Did the player reload instead of switching weapons?
	Scene	Plausible but nonexistent events	Did a vehicle explode in this area?
	Temporal	Real events outside the question context	Did the explosion occur before the firefight?
	Role	Correct event but wrong agent	Did the enemy reload their weapon?
	Intent	Alternative plausible motivations	Did the player reload to prepare for a long fight?
<i>Question Form</i>	Identification	Select the correct answer from options	Which of the following actions occurred?
	Existence	Binary true or false question	Did an explosion occur in this clip?
	Absent	Identify what did not occur	Which event did not happen during healing?
	Intent	Ask why an action was performed	Why did the player reload their weapon?
	Count	Ask for quantities or frequencies	How many enemies appeared in the scene?

Table 7: GameplayQA question taxonomy. Each question is defined by a combination of dimensions, enabling systematic coverage of perception, temporal reasoning, and cross-video reasoning.

Model Name	Version	Inference Provider	Platform
GPT-5	gpt-5	OpenAI	OpenAI
GPT-5-mini	gpt-5-mini	OpenAI	OpenAI
GPT-5-nano	gpt-5-nano	OpenAI	OpenAI
Gemini-2.5-Pro	gemini-2.5-pro	Google	Google AI Studio
Gemini-3-Flash	gemini-3-flash	Google	Google AI Studio
Gemini-2.5-Flash	gemini-2.5-flash	Google	Google AI Studio
Claude-4.5-Sonnet	claude-4.5-sonnet	Amazon Bedrock	OpenRouter
Claude-4.5-Haiku	claude-4.5-haiku	Amazon Bedrock	OpenRouter
Seed-1.6	bytedance-seed/seed-1.6	Seed	OpenRouter
Seed-1.6-Flash	bytedance-seed/seed-1.6-flash	Seed	OpenRouter
Qwen-3-VL-235B	qwen/qwen3-vl-235b-a22b-instruct	Fireworks	OpenRouter
Qwen-3-VL-30B	qwen/qwen3-vl-30b-a3b-instruct	Fireworks	OpenRouter
Qwen-3-VL-8B	qwen/qwen3-vl-8b-instruct	Alibaba	OpenRouter
Gemma-3-27B	google/gemma-3-27b-it	Chutes	OpenRouter
Gemma-3-12B	google/gemma-3-12b-it	Chutes	OpenRouter
Gemma-3-4B	google/gemma-3-4b-it	Chutes	OpenRouter

Table 8: Inference configurations for tested models. All models are evaluated using the listed inference providers.

Model Name	Default Reasoning Mode
GPT-5	Balanced (medium)
GPT-5-mini	Balanced (medium)
GPT-5-nano	Balanced (medium)
Gemini-2.5-Pro	High (dynamic)
Gemini-3-Flash	High (dynamic)
Gemini-2.5-Flash	High (dynamic)
Claude-4.5-Sonnet	Standard (extended thinking off)
Claude-4.5-Haiku	Standard (extended thinking off)

Table 9: Default reasoning effort for tested models. All models are evaluated using their default reasoning settings.

Form	Entity	Code	Template
<i>IDENT</i>	SA	SA-IDENT	Which of the following actions did the POV player perform during the video?
	SS	SS-IDENT	Which of the following best describes the POV player's state in the video?
	OA	OA-IDENT	Which of the following actions did {other} perform during the video?
	OS	OS-IDENT	Which of the following best describes {other}'s state in the video?
	WO	WO-IDENT	Which of the following objects appeared in the video?
	WE	WE-IDENT	Which of the following event occurred in the video?
<i>EXIST</i>	SA	SA-EXIST	Did the POV player perform the action: "{caption}"?
	SS	SS-EXIST	Can you describe the POV player's state as: "{caption}"?
	OA	OA-EXIST	Did the {other} perform the action: "{caption}"?
	OS	OS-EXIST	Can you describe the {other}'s state as: "{caption}"?
	WO	WO-EXIST	Did the object "{caption}" appear in the video?
	WE	WE-EXIST	Did the event "{caption}" occur in the video?
<i>ABSENT</i>	SA	SA-ABSENT	Which action did the POV player NOT perform?
	SS	SS-ABSENT	Which of the following states does not describe the POV player's state?
	OA	OA-ABSENT	Which action did the {other} NOT perform?
	OS	OS-ABSENT	Which of the following does not describe the {other}'s state?
	WO	WO-ABSENT	Which objects is NOT present in the scene?
	WE	WE-ABSENT	Which of the following events did NOT occur in the video?
<i>COUNT</i>	SA	SA-COUNT	How many times did the POV player perform the action: "{caption}"?
	OA	OA-COUNT	How many times did the {other} perform the action: "{caption}"?
	WO	WO-COUNT	How many {caption} are there in the scene?
	WE	WE-COUNT	How many times did the event "{caption}" occur in the video?
<i>INTENT</i>	SA	SA-INTENT	Why did the POV player perform the action: "{caption}"?
	OA	OA-INTENT	Why did the {other} perform the action: "{caption}"?

Table 12: Level 1 (perception) question templates. Entity types: SA (Self-Action), SS (Self-State), OA (Other-Action), OS (Other-State), WO (World-Object), WE (World-Event).

Form	Ref → Ans	Code	Template
<i>IDENT</i>	SA → SS	SA2SS-IDENT	When the POV player was performing the action: “{refCaption}”, which of the following best describes their state?
	SA → OA	SA2OA-IDENT	When the POV player was performing the action: “{refCaption}”, which of the following actions did {other} perform?
	SS → SA	SS2SA-IDENT	When the POV player’s “{refCaption}”, which of the following actions did they perform?
	OA → SS	OA2SS-IDENT	When {other} was performing the action: “{refCaption}”, which of the following best describes the POV player’s state?
	WO → SA	WO2SA-IDENT	At the moment when the object “{refCaption}” appeared, which of the following actions did the POV player perform?
	WE → OA	WE2OA-IDENT	At the moment when the event “{refCaption}” occurred, which of the following actions did {other} perform?
<i>EXIST</i>	SA → SS	SA2SS-EXIST	When the POV player was performing the action: “{refCaption}”, can you describe their state as: “{caption}”?
	SA → OA	SA2OA-EXIST	When the POV player was performing the action: “{refCaption}”, did {other} perform the action: “{caption}”?
	SS → SA	SS2SA-EXIST	When the POV player’s “{refCaption}”, did they perform the action: “{caption}”?
	OA → SS	OA2SS-EXIST	When {other} was performing the action: “{refCaption}”, can you describe the POV player’s state as: “{caption}”?
	WO → SA	WO2SA-EXIST	At the moment when the object “{refCaption}” appeared, did the POV player perform the action: “{caption}”?
	WE → OA	WE2OA-EXIST	At the moment when the event “{refCaption}” occurred, did {other} perform the action: “{caption}”?
<i>ABSENT</i>	SA → SS	SA2SS-ABSENT	When the POV player was performing the action: “{refCaption}”, which of the following does NOT describe their state?
	SA → OA	SA2OA-ABSENT	When the POV player was performing the action: “{refCaption}”, which action did {other} NOT perform?
	SS → SA	SS2SA-ABSENT	When the POV player’s “{refCaption}”, which action did they NOT perform?
	OA → SS	OA2SS-ABSENT	When {other} was performing the action: “{refCaption}”, which of the following does NOT describe the POV player’s state?
	WO → SA	WO2SA-ABSENT	At the moment when the object “{refCaption}” appeared, which action did the POV player NOT perform?
	WE → OA	WE2OA-ABSENT	At the moment when the event “{refCaption}” occurred, which action did {other} NOT perform?

Table 13: Level 2 entity-reference question templates (representative examples; full set covers all 30 Ref → Ans pairs across SA, SS, OA, OS, WO, WE). Code format: {Ref}2{Ans}-{Form}. Each EXIST question is further instantiated with distractor subtypes: *True*, *Lexical*, *Scene*, *Temporal*, and *Role*.

Form	Ans	Code	Template
<i>IDENT</i>	SA	TR2SA-IDENT	During {timestamp}, which of the following actions did the POV player perform?
	SS	TR2SS-IDENT	During {timestamp}, which of the following best describes the POV player's state?
	OA	TR2OA-IDENT	During {timestamp}, which of the following actions did {other} perform?
	OS	TR2OS-IDENT	During {timestamp}, which of the following best describes {other}'s state?
	WO	TR2WO-IDENT	During {timestamp}, which of the following objects appeared?
	WE	TR2WE-IDENT	During {timestamp}, which of the following events occurred?
<i>EXIST</i>	SA	TR2SA-EXIST	During {timestamp}, did the POV player perform the action: "{caption}"?
	SS	TR2SS-EXIST	During {timestamp}, can you describe the POV player's state as: "{caption}"?
	OA	TR2OA-EXIST	During {timestamp}, did {other} perform the action: "{caption}"?
	OS	TR2OS-EXIST	During {timestamp}, can you describe {other}'s state as: "{caption}"?
	WO	TR2WO-EXIST	During {timestamp}, did the object "{caption}" appear?
	WE	TR2WE-EXIST	During {timestamp}, did the event "{caption}" occur?
<i>ABSENT</i>	SA	TR2SA-ABSENT	During {timestamp}, which action did the POV player NOT perform?
	SS	TR2SS-ABSENT	During {timestamp}, which of the following does NOT describe the POV player's state?
	OA	TR2OA-ABSENT	During {timestamp}, which action did {other} NOT perform?
	OS	TR2OS-ABSENT	During {timestamp}, which of the following does NOT describe {other}'s state?
	WO	TR2WO-ABSENT	During {timestamp}, which object did NOT appear?
	WE	TR2WE-ABSENT	During {timestamp}, which of the following events did NOT occur?

Table 14: Level 2 timestamp-reference (TR) question templates. The placeholder {timestamp} is replaced by a formatted time range such as [00:01 to 00:12]. Code format: TR2{Ans}-{Form}.

Type	Ref → Ans	Code	Template
<i>Cross-Video Reference (V1-$\{Ref\}$2V2-$\{Ans\}$-$\{Form\}$)</i>			
<i>IDENT</i>	SA → SA	V1-SA2V2-SA-IDENT	When POV1 player was performing the action: “ $\{refCaption\}$ ”, which of the following actions did POV2 player perform at the same time?
	SA → SS	V1-SA2V2-SS-IDENT	When POV1 player was performing the action: “ $\{refCaption\}$ ”, which of the following best describes POV2 player’s state at the same time?
	OA → SA	V1-OA2V2-SA-IDENT	When $\{refOther\}$ was performing the action: “ $\{refCaption\}$ ” in POV1, which of the following actions did POV2 player perform at the same time?
	WE → WO	V1-WE2V2-WO-IDENT	At the moment when the event “ $\{refCaption\}$ ” occurred in POV1, which of the following objects appeared in POV2 at the same time?
<i>EXIST</i>	SA → SA	V1-SA2V2-SA-EXIST	When POV1 player was performing the action: “ $\{refCaption\}$ ”, did POV2 player perform the action: “ $\{caption\}$ ” at the same time?
	SA → OA	V1-SA2V2-OA-EXIST	When POV1 player was performing the action: “ $\{refCaption\}$ ”, did $\{other\}$ perform the action: “ $\{caption\}$ ” in POV2 at the same time?
	OS → WE	V1-OS2V2-WE-EXIST	When $\{refOther\}$ ’s “ $\{refCaption\}$ ” in POV1, did the event “ $\{caption\}$ ” occur in POV2 at the same time?
	WO → SS	V1-WO2V2-SS-EXIST	When the object “ $\{refCaption\}$ ” appeared in POV1, was POV2 player’s “ $\{caption\}$ ” at the same time?
<i>POV Identity (POV-ID): Which video corresponds to the player who did X?</i>			
<i>POV-ID</i>	SA	SA-POV-ID	Which video corresponds to the player who performed the action: “ $\{caption\}$ ”?
	SS	SS-POV-ID	Which video corresponds to the player whose “ $\{caption\}$ ”?
	OA	OA-POV-ID	Which video shows $\{other\}$ performing the action: “ $\{caption\}$ ”?
	OS	OS-POV-ID	Which video shows $\{other\}$ whose “ $\{caption\}$ ”?
	WO	WO-POV-ID	Which video shows the object “ $\{caption\}$ ”?
	WE	WE-POV-ID	Which video shows the event “ $\{caption\}$ ”?
<i>Temporal Ordering (ORDER): Which happened first?</i>			
<i>ORDER</i>	SA	SA-ORDER	Which of the following actions happened first?

Table 15: Level 3 (cross-video) question templates. Cross-video reference templates cover all 36 Ref → Ans pairs (6 reference types × 6 answer types) in both IDENT and EXIST forms; representative examples are shown. V1/V2 are replaced by actual video indices at generation time. $\{refOther\}$ refers to the other player in the reference video. POV-ID answer options are video numbers; ORDER options are formatted as “The POV player in Video X is [action]”.